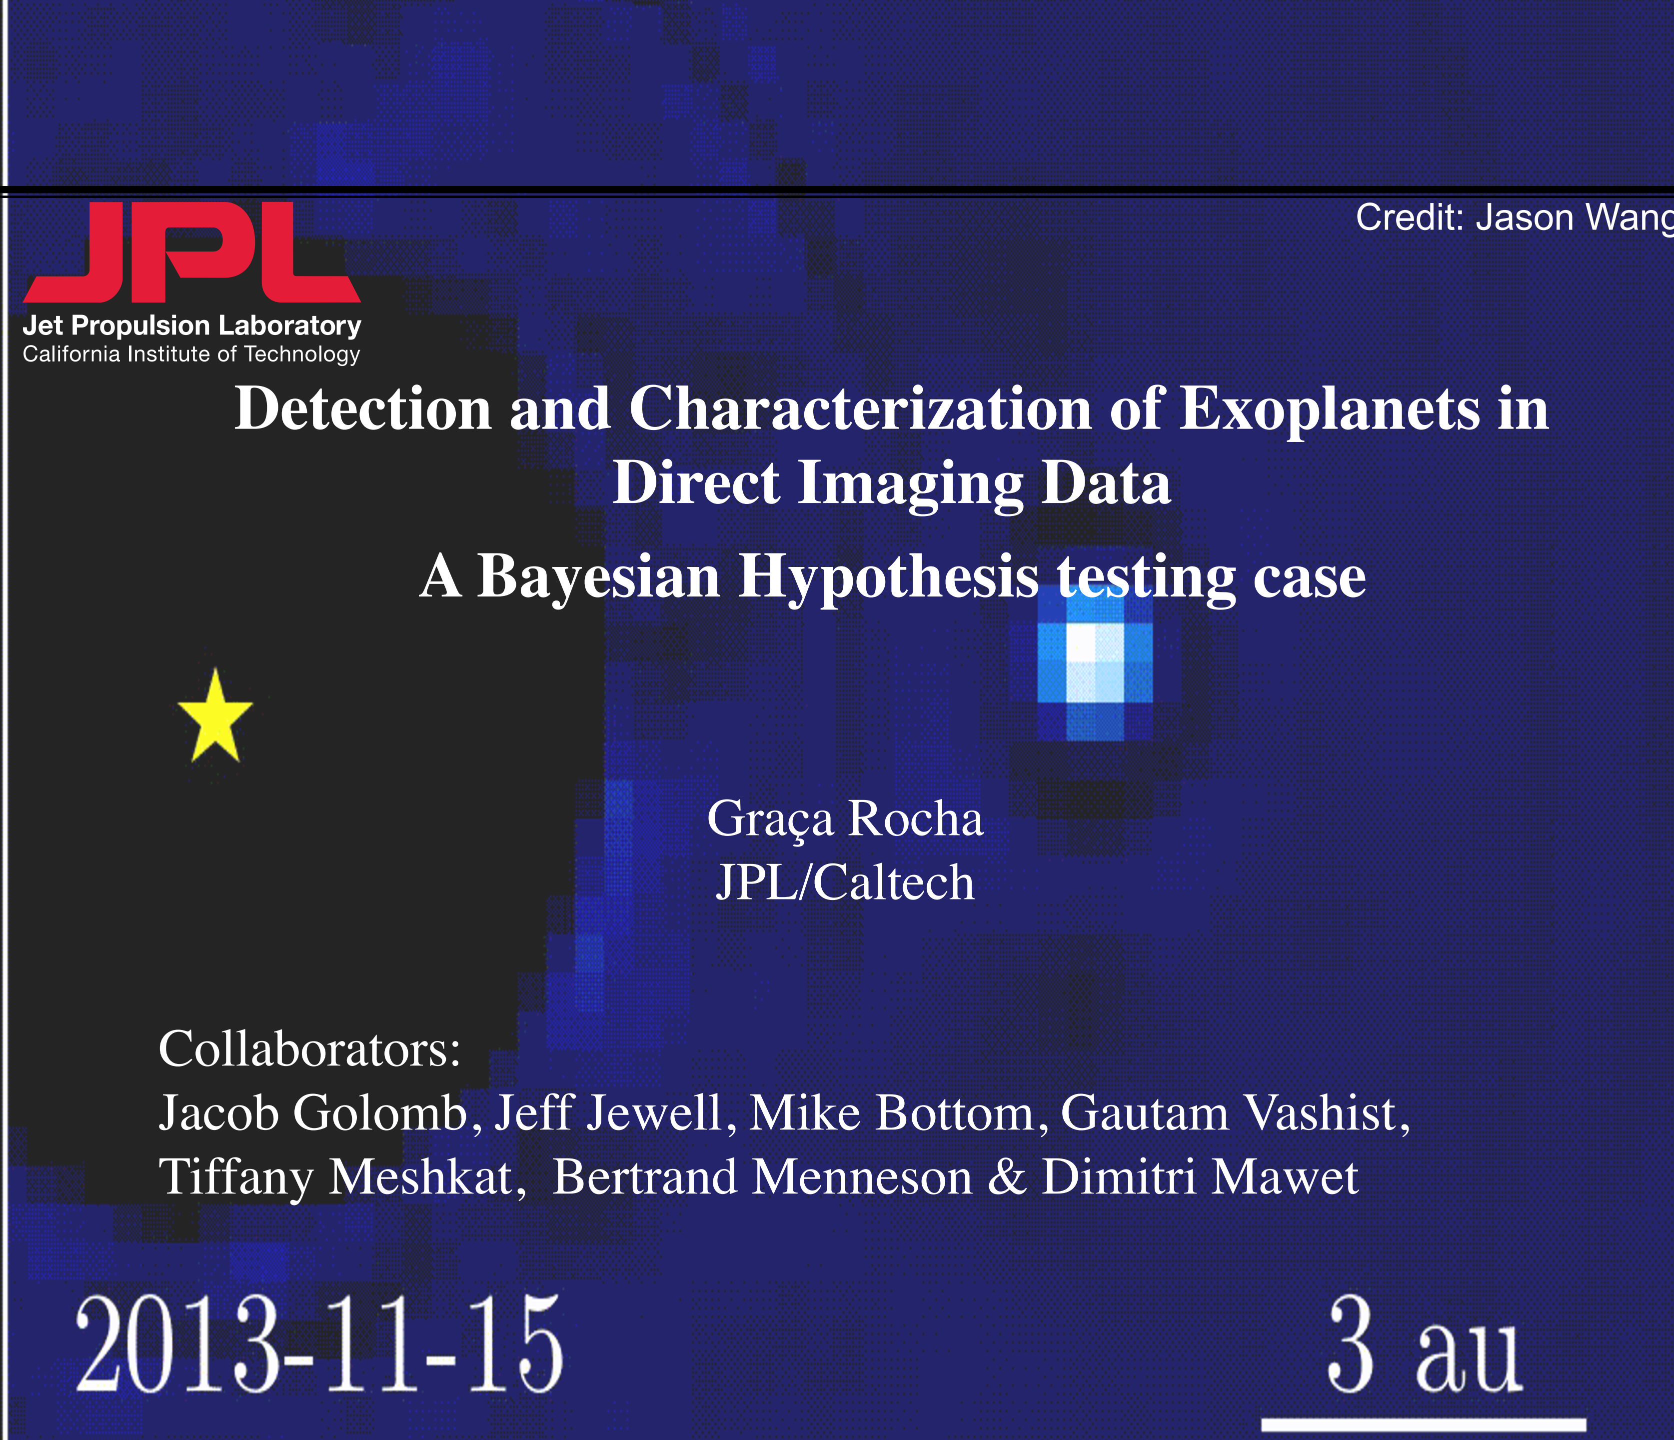

Jet Propulsion Laboratory
California Institute of Technology

Detection and Characterization of Exoplanets in Direct Imaging Data

A Bayesian Hypothesis testing case

Graça Rocha
JPL/Caltech

Collaborators:
Jacob Golomb, Jeff Jewell, Mike Bottom, Gautam Vashist, Tiffany Meshkat, Bertrand Menneson & Dimitri Mawet

I - Synergies of CMB and Exoplanets Data Analysis

❑ Step I:

- Identify the synergies between:

1. Detection + characterization of point sources and SZ clusters in CMB data
2. Detection + characterization of exoplanets in direct imaging data

➢ For both cases there is the “noise” or “systematics” component that can mimic the signal: CMB and speckles

- In 1 –> background:

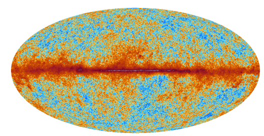

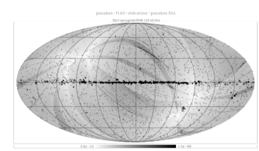

Point Sources

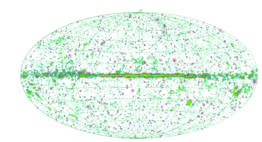

  - CMB (correlated noise), instrumental noise (assumed white) and diffuse foreground emission (Non-Gaussian, ignored here as it is treated separately),..

- In 2 –> background:

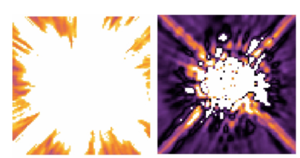

Exoplanets

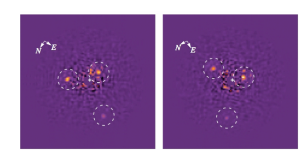

  - Speckles (stellar PSF), instrumental noise (assumed white),...

• Step II:

- Phrase the detection and characterization within a Bayesian framework

Bayesian
What does it really mean?

- What does it mean to recast the problem of planet detection and characterization into a Bayesian perspective?

A. The Bayesian framework entails defining the following key ingredients:

a data model + *a Likelihood shape* + *model parameter priors*

$\mathcal{D}(x)=s(x)+n(x)$ $\mathrm{L}(d)=\mathcal{P}(d|\Theta,\mathcal{H})$ $\Pi(\Theta)=\mathcal{P}(\Theta|\mathcal{H})$

B. Next apply Bayes Theorem – to retrieve the probability distribution of the model parameters:

Bayes Theorem $\longrightarrow$ *Posterior distributions of the model* + *Best Fit models*

$\mathcal{P}(\Theta|d,\mathcal{H})$

eg. Maximum Likelihood

Inference → Parameter estimation → $\mathcal{P}(\Theta|d,\mathcal{H}) = \mathcal{P}(d|\Theta,\mathcal{H})\ \mathcal{P}(\Theta|\mathcal{H}) / \mathcal{P}(d|\mathcal{H})$

Ignored

Model Selection → Evidence is crucial $\mathcal{E}=\mathcal{P}(d|\mathcal{H}) = \mathcal{Z}$

Expectation of the likelihood over the prior → $\mathcal{Z} = \int L(\Theta)\pi(\Theta)d^D\Theta,$

Bayesian Inference
Basic tools

- In contrast to parameter estimation problems → in model selection the evidence takes central role and is simply the factor required to normalize the posterior:

Evidence

$$\mathcal{Z} = \int L(\Theta)\pi(\Theta)d^D\Theta,$$

Evaluation of this multidimensional Integral is a challenging numerical task – resort to sampling techniques: MCMC, Multinest, (Sivia &Skilling 2006; Feroz et al. 2009), etc. or model the posterior as a multivariate Gaussian centered at its peak(s) and apply the Laplace formula (Hobson, Bridle & Lahav 2002).

- The evidence automatically implements Occam's razor:

A simpler theory with compact parameter space will have a larger evidence than a more complicated one, unless the latter is significantly better at explaining the data.

- Model selection between two models H_0 and H_1 can be decided by comparing their respective posterior probabilities given the observed data set d:

$$\frac{\Pr(H_1|d)}{\Pr(H_0|d)} = \frac{\Pr(d|H_1)\Pr(H_1)}{\Pr(d|H_0)\Pr(H_0)} = \frac{\mathcal{Z}_1}{\mathcal{Z}_0}\frac{\Pr(H_1)}{\Pr(H_0)},$$

$\Pr(H_1)/\Pr(H_0)$ = prior probability ratio for the models

Direct Imaging – background

- In direct imaging data, the image is made up of stellar PSF, noise, and an astronomical signal, such as a planet in the image

- The goal of most of current methodologies for processing these images involves modelling the PSF/residual background, subtracting it from the target image, and looking for bright point sources as planets

  - The search/detection is conducted within a frequentist perspective
    - A *threshold* is set a priori

- This work
  - The search/detection is conducted within a Bayesian framework
    - Bayesian decision theory – *Evidence* based Model Selection

JPL

An example: KLIP

Target image (T) is made up of the background, Point Spread Function (PSF) from the star, and the astronomical source (A).

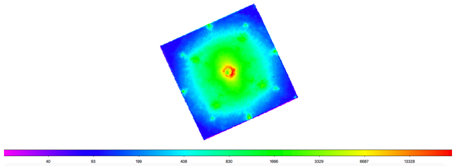

$$T(n) = I_{\Psi_0}(n) + \epsilon A(n)$$

Epsilon is either 0 or 1, depending on whether or not there is the astrophysical source in the pixel (indexed by n). Note these are all 1-D arrays.

Goal: Model and subtract it from the target image to get the planet.

The way we construct the stellar PSF models is
by using Principal Component Analysis (PCA)
Ie create a new basis that correspond to the
most variability (spread) amongst correlated points

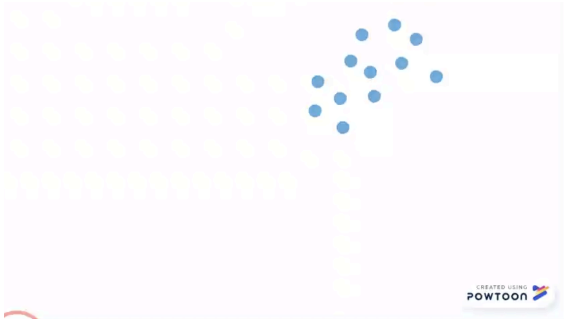

Graphically, we then end up with new axes that form a new basis that makes these points uncorrelated

An example: KLIP

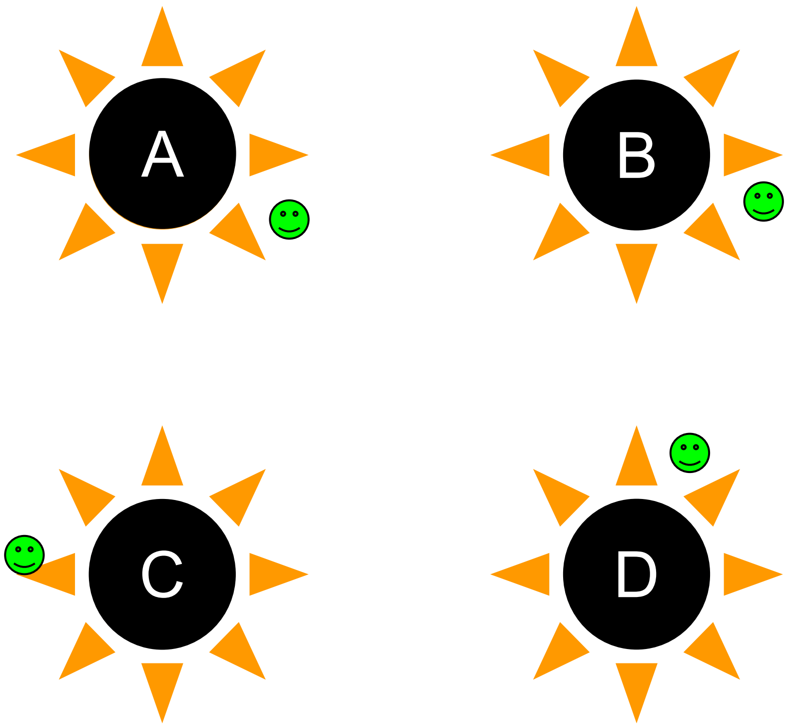

Choosing the reference images involves choosing the set of images where any potential planet in the image would have moved the least amount, so that the variation seen in the images is due to the stellar PSF. This is calculated based on the star's coordinates and the parallactic angle.

Black circles represent the star, blocked by the coronagraph. The orange is the residual stellar PSF. The green is a particular location in the image, corresponding to a test location around the star which has moved with respect to the star due to the sky rotating in the frame over time.

Constructing the KLIP

Stack the reference images in time, and compute the covariances of the pixels:

$$Cov[i,j] = \sum_{n=1}^{N} R_i(n) R_j(n)$$

We now change dimensionality to construct an orthogonal basis whose vectors show the pattern of greatest variations, in descending order. These KL modes are constructed by:

$$Z_k^{KL} = \frac{1}{\sqrt{\Lambda_k}} \sum_{p=1}^{K} C_k(\Psi_p) R_p(n)$$

Where C_K are the eigenvectors of the covariance matrix and R is a reference image. The resulting $Z^{KL}{}_k$ vectors are orthogonal and characterize the modes of variability.

Here we use PyKLIP to get the KL modes: https://pyklip.readthedocs.io/en/latest/index.html

Say the blue dots are pixel values and Z_1 and Z_2 are the orthogonal KL mode vectors

Projecting the blue dots (pixels) onto each axis tells us how much variability lies along that mode, or how much variation lies along those sets of pixels

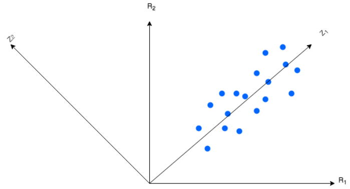

Model of stellar PSF

Finally, we get a best-estimate model for the stellar PSF:

$$I_{\Psi_0}(n) = \sum_{k=1}^{K_{klip}} < T, Z_k^{KL} > Z_k^{KL}(n)$$

Where K_{KLIP} is some limit where we cut off the KL modes. A statistical method for determining this limit would be future work.

The inner product above is simply a scalar characterizing the magnitude of the variability in that mode.

Assuming K_{KLIP} is the optimal number, we now have a full mathematical description of the stellar PSF.

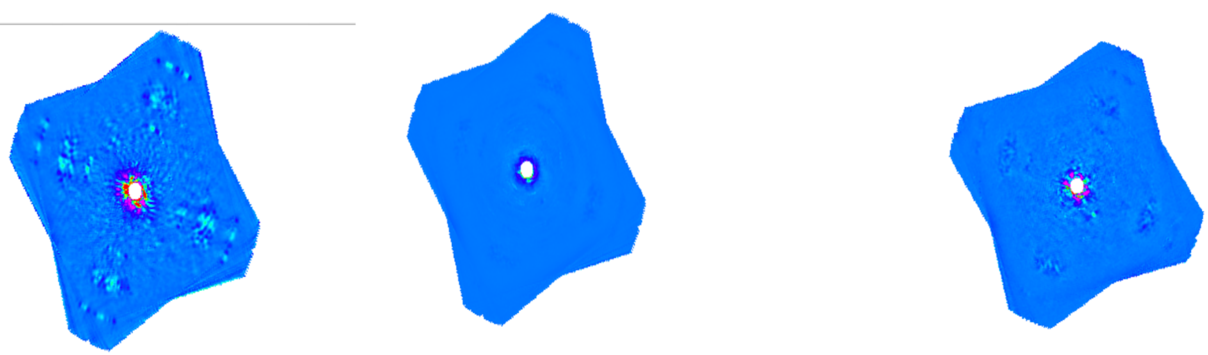

Using Gemini Planetary Imager Data

Satellite spots are manually created in the optics, the locations of which on the CCD are mentioned in the FITS header.

The satellite spots for the image are averaged across all frames for that wavelength. Since the satellite spots are formed as point sources, this serves as a model of the instrumental PSF.

Note that the PSFs are NOT modelled as simply Gaussian approximations.

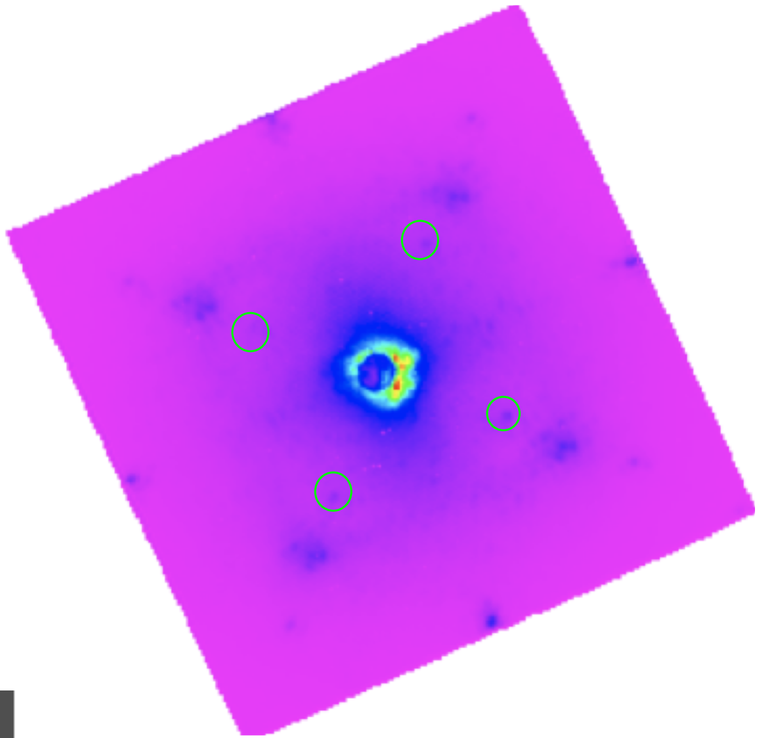

Locations of satellite spot PSFs circled in green

Actual GPI instrumental PSF →

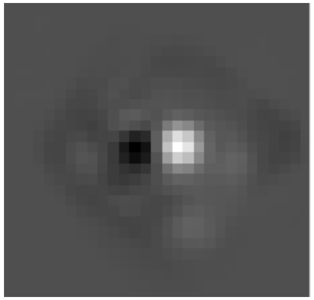

JPL

What we have now:

- A noise-subtracted image, with a location of a potential planet in mind
- A model of the instrumental PSF, which will reflect how the planet will appear in the image

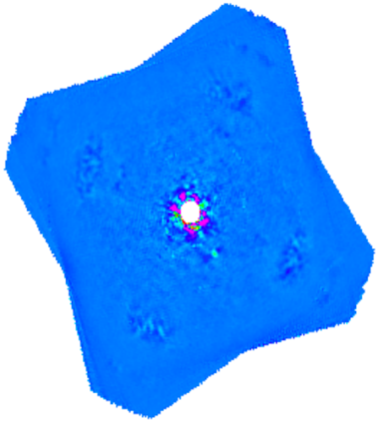

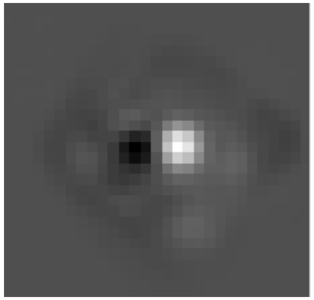

- We have a guessed location for the planet
- We use forward modeling to get more accurate astrometry using a model of the planet - Essentially, forward modeling tackles the question:

  How would a planet appear if it were truly in the image?

Bayesian Detection & Characterization

- Create single image from mean-combining temporal and frequency data slices
- Construct the data model for the co-added data
- Construct the Likelihood using approach followed in KLIP-FM
  - Model the likelihood as a Gaussian of the residuals R :

$$R(x_p, y_p, \alpha) = (D - \alpha F(x_p, y_p))_{\mathcal{F}}$$

D = data, planet model = α F(x_p,y_p), F = fitting region(a rectangle fixed on the location of the planet in the data). The flux, F, is the modelled instrumental PSF from the forward modeling, scaled to the data with the scale factor α

$$ln\mathcal{L} = -\frac{1}{2}(R^T C^{-1} R + ln(detC) + N_{pix} ln(2\pi))$$

- For the residual speckles, we must consider *correlated* background

$$C_{ij} = \sigma_i \sigma_j (1 + \frac{\sqrt{3} r_{ij}}{\ell}) e^{\frac{-\sqrt{3} r_{ij}}{\ell}} \qquad \ell \approx \frac{\lambda}{D} \approx 3$$

- Where C is the *Matern* covariance function: (Wang et al. 2016)

- Apply *Nested sampling* to get evidence of competing models and posterior distributions of the model parameters

JPL

•

Signal or noise?

- We will show that it is possible to detect dim sources
- This now moves onto the problem of *characterization*
- We have detected *something* that minimizes the residuals, but is it more likely to be a planet or still background noise?

To do this, we require 2 models

We run two models on the same location and compare the evidences.

The hypothesis (planet) model uses the forward model with α as a parameter, and corresponds to a planet present in the image.

The null hypothesis is $\alpha = 0$ and thus means a planet is not present within the prior bounds.

With strong enough evidence for the planet model as compared to the null hypothesis, the null hypothesis can be rejected and it can be claimed with a degree of confidence that a planet is present.

What do you mean "compare the evidences of the models"?

Good question!

H_1: Planet Model

$$R(x_p, y_p, \alpha) = (D - \alpha F(x_p, y_p))_{\mathcal{F}}$$

$$ln\mathcal{L} = -\frac{1}{2}(R^T C^{-1} R + ln(detC) + N_{pix} ln(2\pi))$$

We compute the evidence as:

$$Z_1 = \int_{\theta=\ell,x,y,\alpha} \mathcal{L}(D|\theta)\pi(\ell)\pi(x)\pi(y)\pi(\alpha)d\theta$$

H_0: No Planet Present

- i.e. model the correlated noise

$$ln\mathcal{L} = -\frac{1}{2}(D^T C^{-1} D + ln(detC) + N_{pix} ln(2\pi))$$

- Note that here, our likelihood is simply the likelihood of *just the data* with the Matern covariance matrix
  - In effect, $\alpha = 0$
- Evidence:

$$Z_0 = \int_{\theta=\ell,x,y} \mathcal{L}(D|\theta)\pi(\ell)\pi(x)\pi(y)d\theta$$

How do we compute this evidence term?

- To compute the evidence term, we use nested sampling (Skilling 2004)
- Nested sampling turns a multidimensional space into a likelihood function of the probability getting that likelihood or greater, thus transforming this multidimensional integral into an integral over one dimension
- Drops N samples in prior space, computes the corresponding likelihoods, and banks the lowest likelihood value

Nested Sampling in PyKLIP

- PyMultiNest is a python wrapper for a Fortran code that does nested sampling given a likelihood model and priors (PyMultiNest: Johannes Buchner, MultiNest: Feroz et al 2008)

- I implemented this nested sampling adaptation into the PyKLIP forward modeling, so we get both the marginals for each parameter as well as the overall evidence

- We can then compare the evidence terms for both the H_0 and H_1 locations run on the same location

Location of beta Pic b

H_1

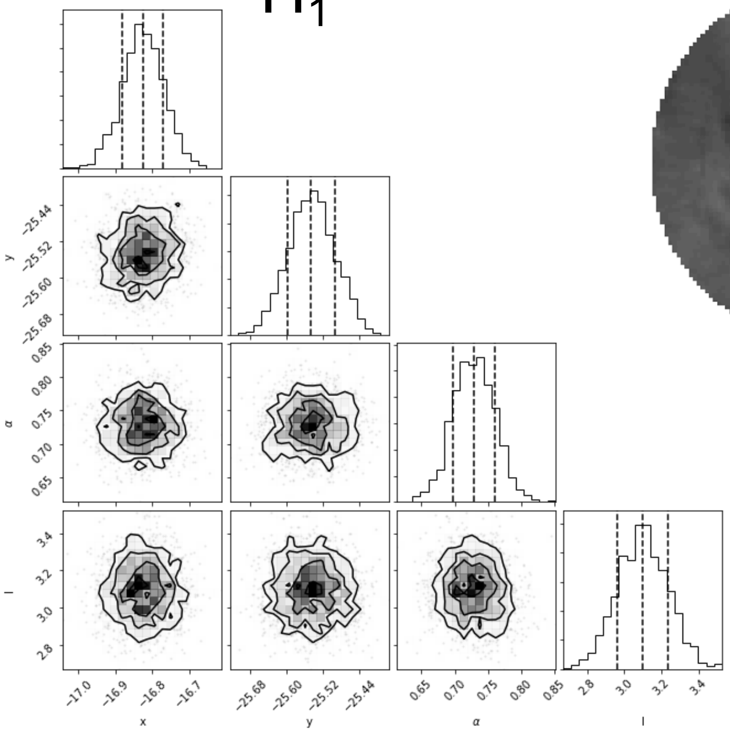

Evidence: lnZ = -463

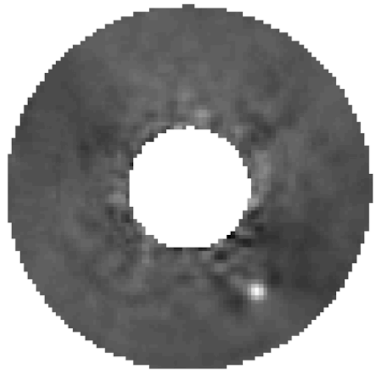

H_0

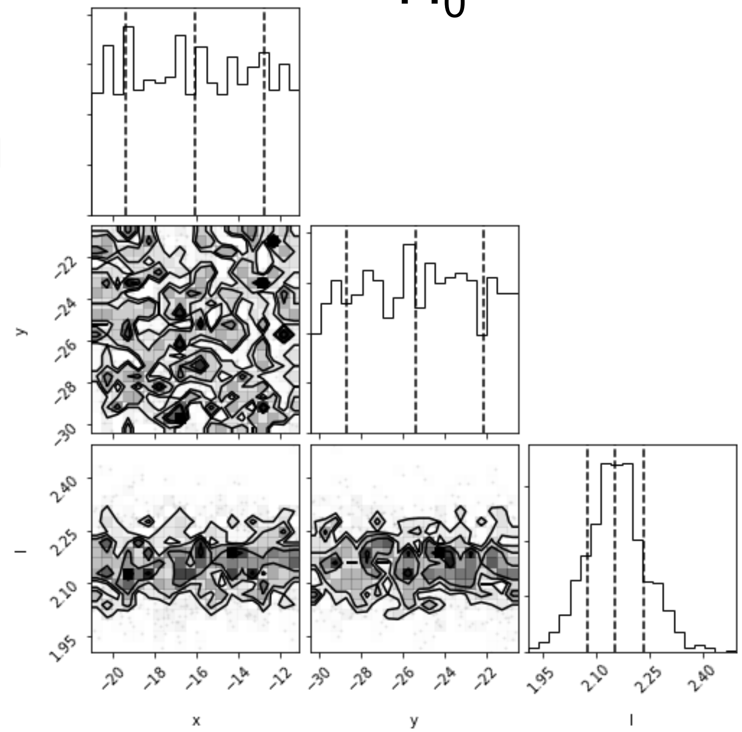

Evidence: lnZ = -638

Model Comparison

Evidence ratios can tell us which model is favored by the data.

$$Ratio = \frac{prob(D|H_1)}{prob(D|H_0)} = \frac{Z_1}{Z_0}$$

Where H_0 is the null hypothesis ($\alpha = 0$, no planet) and H_1 is the "planet" model with α as a parameter.

In this case: $$Ratio = \frac{e^{-463}}{e^{-638}} = 2 * 10^{76}$$

We can reject the null hypothesis with confidence. In other words, we are able to properly quantify the confidence with which we distinguish beta Pic b from the noise/residual speckles.

Synthetic planet at 25% FM Flux

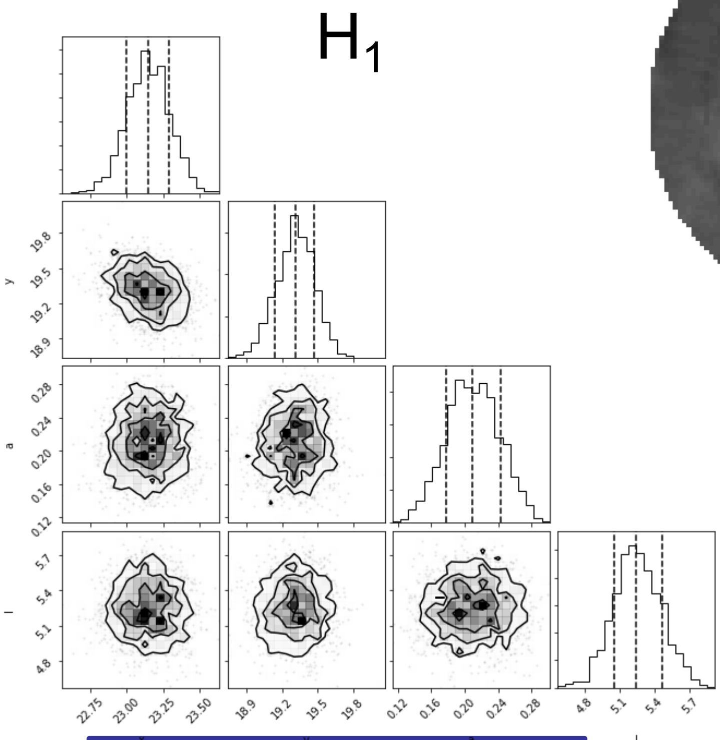

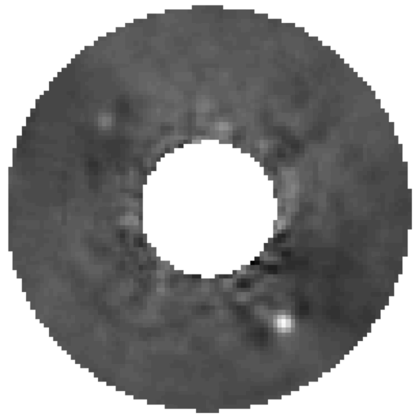

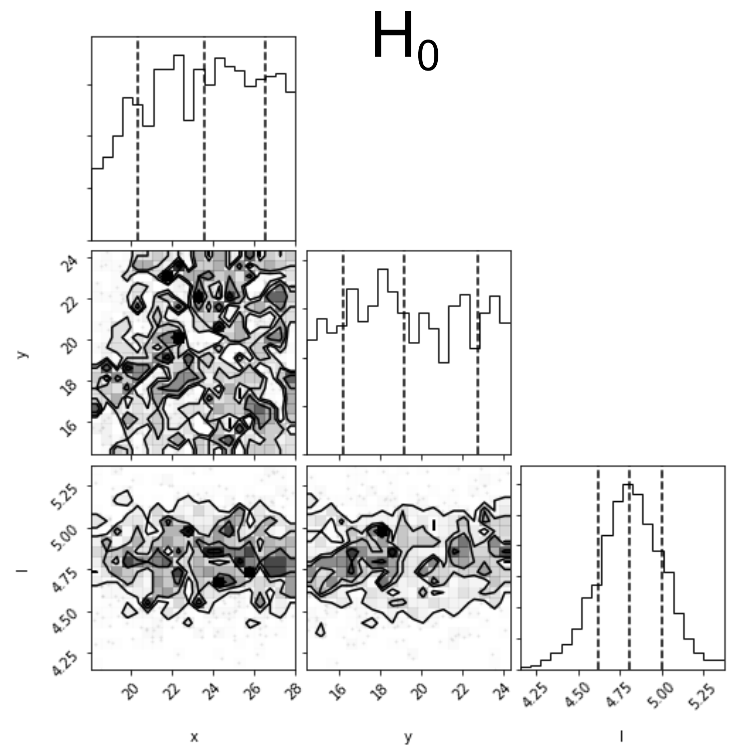

Evidence: $\ln Z = -385$

Ratio = 10^6

SNR from residuals: 6.7σ

SNR from annulus: 3.8σ

Evidence: $\ln Z = -399$

Synthetic planet at 15% FM Flux

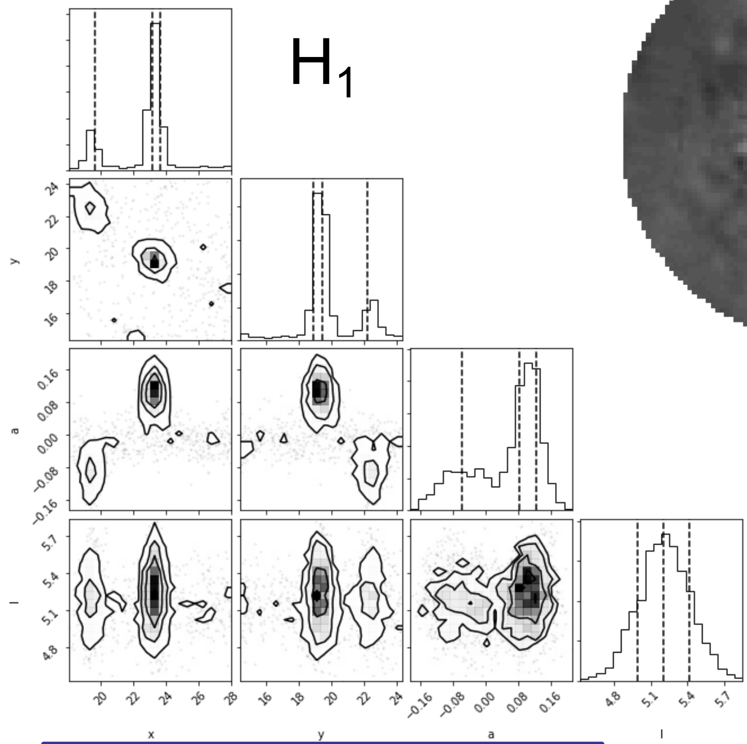

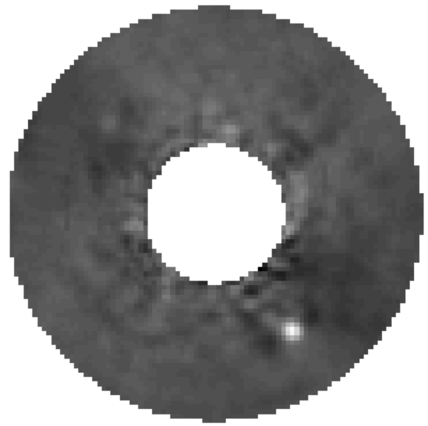

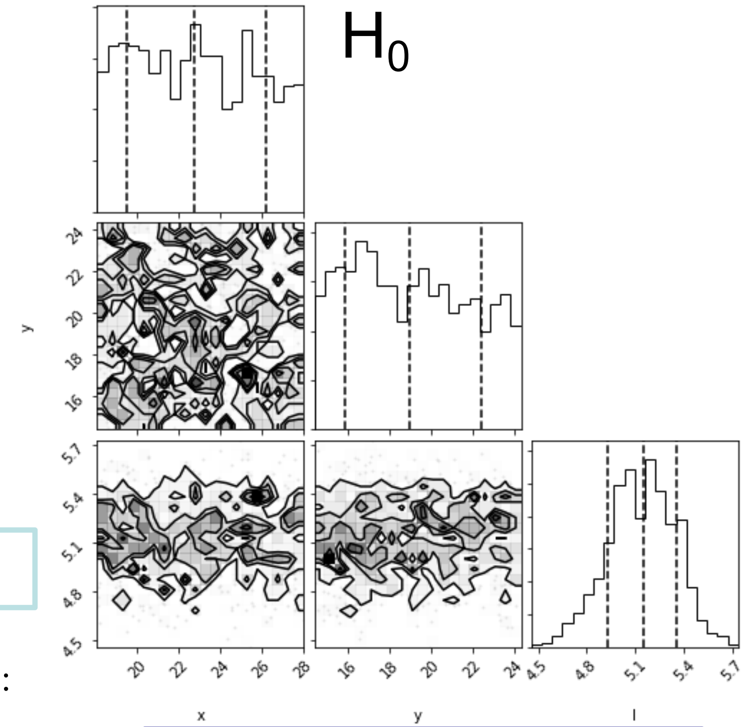

Ratio = 0.29

SNR from residuals: 3.3σ
SNR from annulus: 2.3σ

Evidence: lnZ = -384.9

Evidence: lnZ = -383.67

Synthetic Planet at 15% FM Flux in Different Position

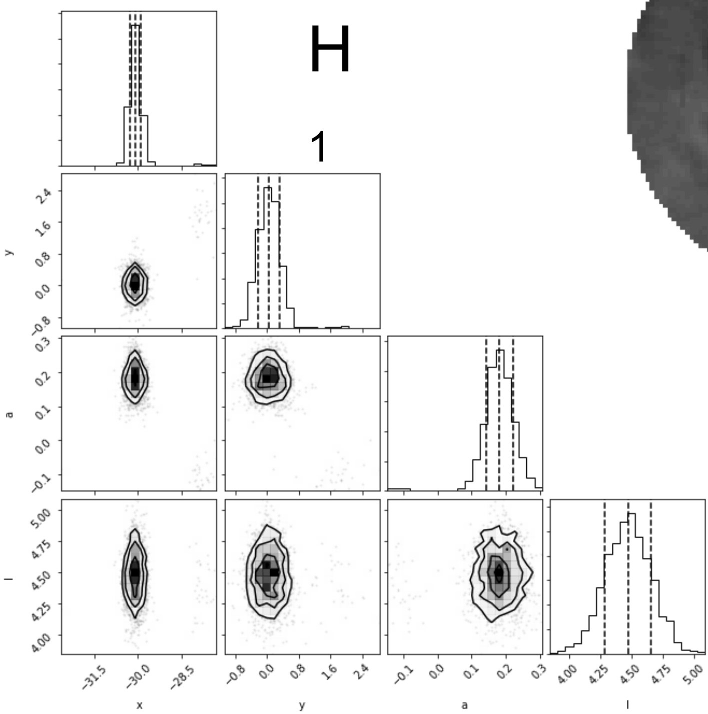

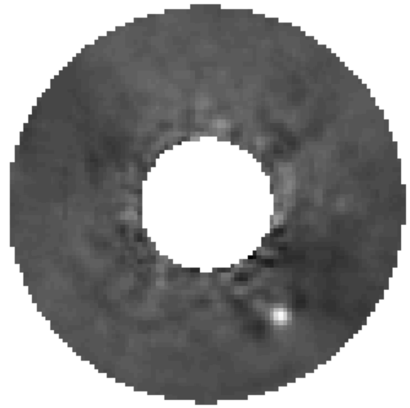

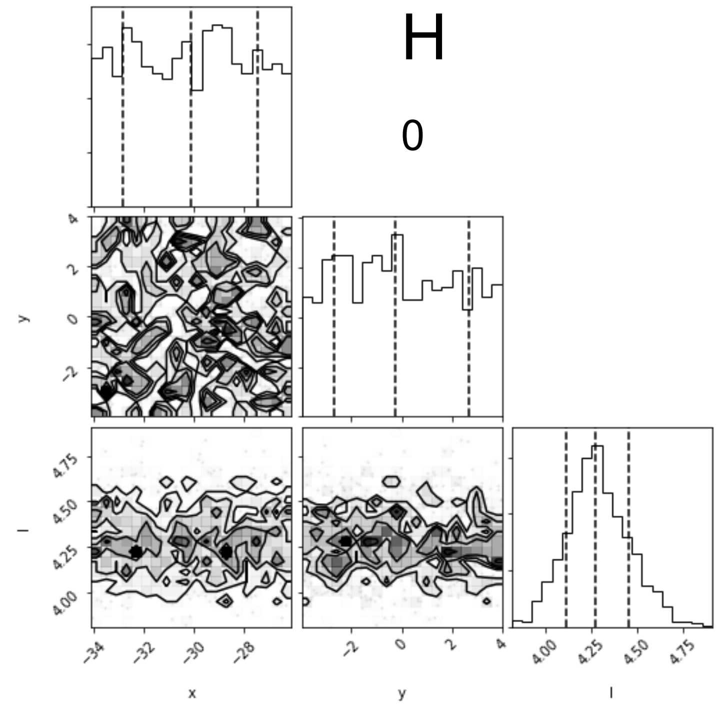

Ratio = 20

SNR from residuals: 3.8σ
SNR from annulus: 2.4σ

Evidence: lnZ = -415

Evidence: lnZ = -418

Synthetic Planet at 10% FM Flux at this Position

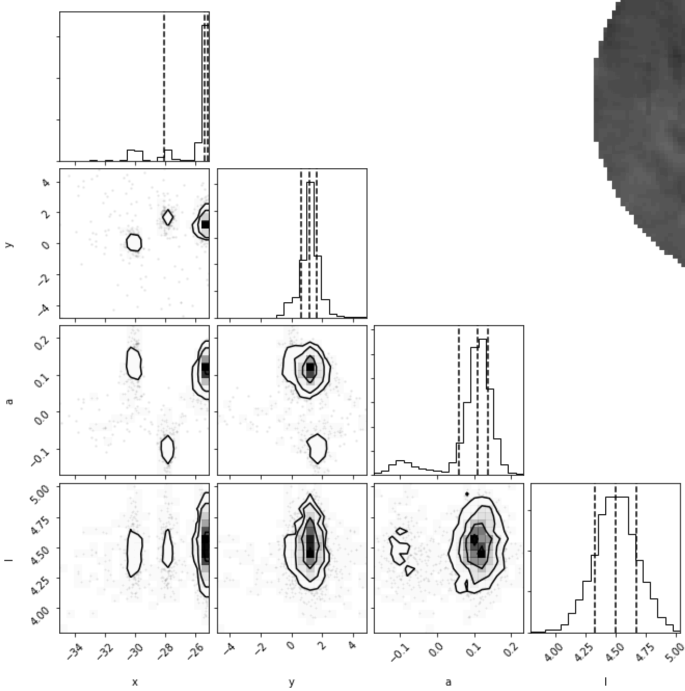

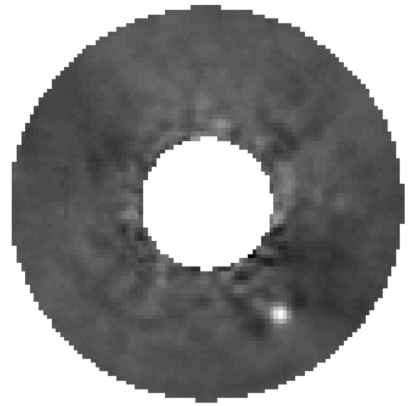

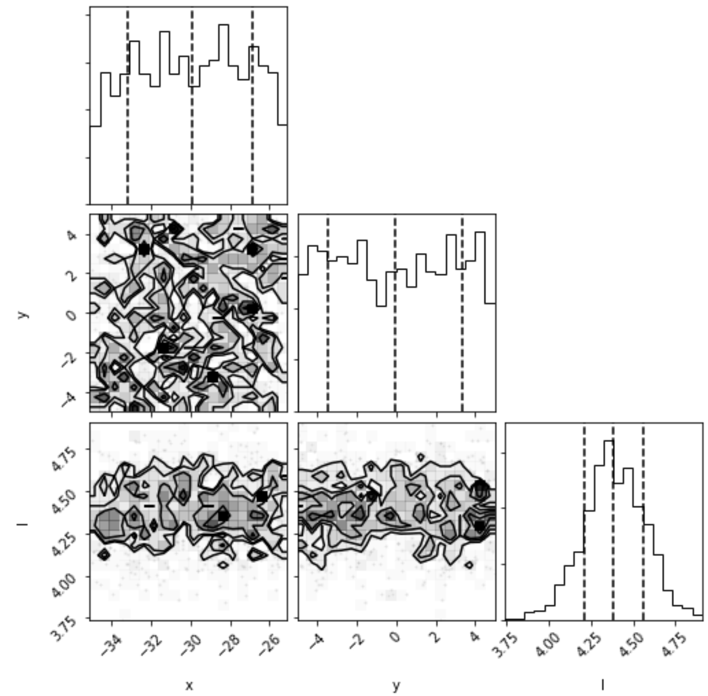

Evidence: lnZ = -413

Ratio = 1.56

SNR from residuals: 2.4σ
SNR from annulus: 1.6σ

Evidence: lnZ = -413.45

Running on noise

H_1

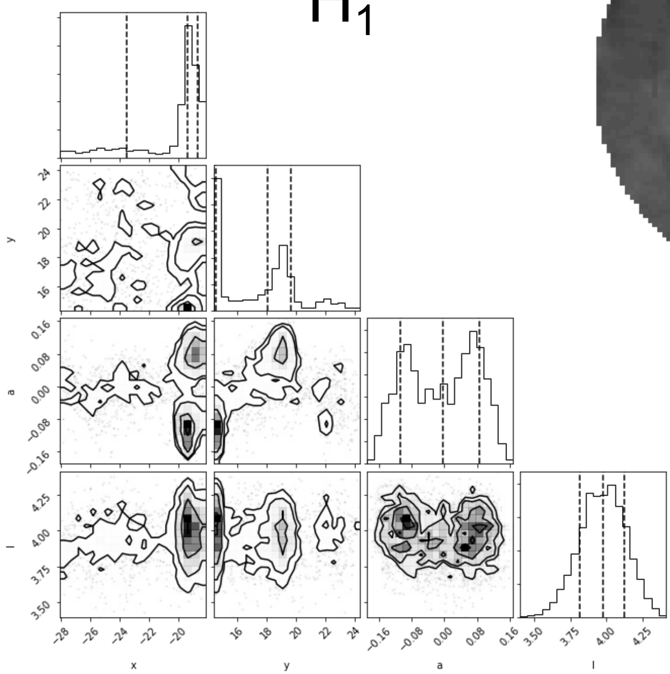

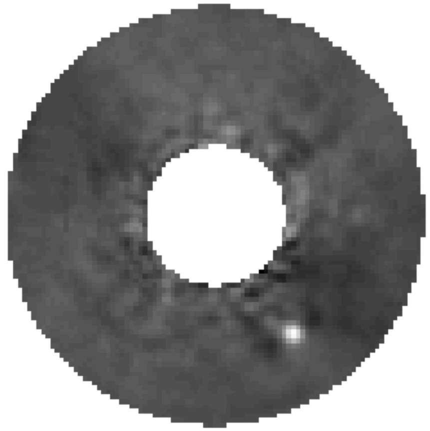

H_0

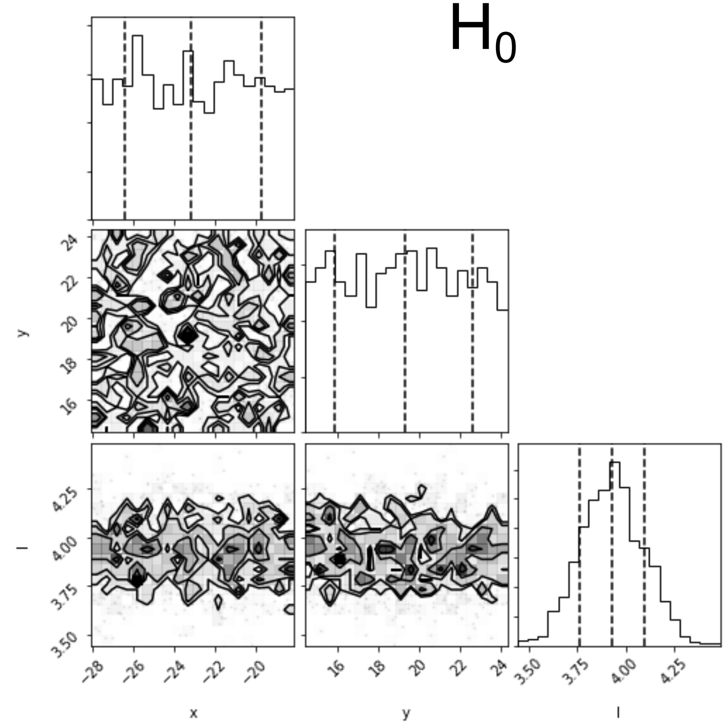

Ratio = 0.19

Evidence: lnZ = −446.26

Evidence: lnZ = -444.6

Conclusions

- Using forward modeling, faint potential planets can be detected
- Nested sampling can be used successfully to distinguish between the planet and background noise
  - The confidence for this characterization depends on the noise level in that region as well as the brightness of the source
- Future: Bayesian Blind Detection— By forward modeling and running the characterization functions on locations of interest in the image, a fully-automated algorithm may be able to detect faint planets in a direct image and characterize them as such.